



DFS Changes in Windows Server 2008, Improvements and Best Practices

DFS has undergone incremental changes, all offering a basic solution for file management across LANs and WANs but the story is incomplete. In this E-Guide, brought to you by SearchWindowsServer.com and Peer Software, learn what improvements have been made to DFS in Windows Server 2008 R2 and what problems still exist. Discover changes that have been made in Windows Server 2008 that make the file system virtualization process a bit less confusing for end users.

Sponsored By:





DFS Changes in Windows Server 2008, Improvements and Best Practices

Table of Contents:

[Key DFS improvements in Windows Server 2008 R2](#)

[New use cases: File System Virtualization](#)

[Using DFS to create file system virtualization in Windows Server 2008](#)

[DFS – An Incomplete Solution](#)

[Resources from Peer Software](#)

Key DFS improvements in Windows Server 2008 R2

By Brien M. Posey, Contributor

Distributed File System (DFS) has been part of Windows Server for many years, and while it has matured with time, certain scalability problems have remained.

For example, while DFS can be scaled to work in large environments, you'll find that configuring, managing and troubleshooting the system becomes exponentially more complicated as the size of the deployment increases.

Fortunately, Microsoft has addressed these scalability and troubleshooting problems with Windows Server 2008 R2.

Scalability enhancements

For the most part, the scalability improvements made to DFS in Windows Server 2008 R2 are best suited for organizations with multiple branch offices

These types of organizations often have a hub and spoke topology, which means their DFS servers are located in a main office – the hub – with the servers' contents replicated to smaller DFS servers in branch offices.

The problem with this type of architecture is that if a DFS server in the main office fails catastrophically, then all of the branch offices could be impacted.

While creating additional replicas in the main office would be one solution, depending on the replication topology's configuration, these additional replicas may not be able to push updates to the branch offices.

Another solution would be a full mesh topology; however, this path is often avoided because of the expense of the extra WAN links and the volume of replication traffic.

Ultimately, the answer for many organizations is to create a replication with two hub members. With this in mind, it's obvious why one of the most welcome new features in R2 is DFS support for failover clusters. Basically, clustering the hub servers in the main office can prevent the branch office replicas from becoming cut off by a hub server failure.

Another improvement in DFS with Windows 2008 R2 is the ability to create read-only replicated folders.

In the past, if you needed a user at a branch office to access data in a replicated folder -- but you didn't want them to change the data -- you needed to use access control lists (ACL) to grant the particular user read-only access. This requires a lot of administrative effort, especially if the branch offices have a high employee turnover.

A new alternative is to create DFS replicas in the branch offices that contain read-only replicated folders, which has the same effect as granting users read-only access to a traditional replicated folder.

New troubleshooting capabilities

In the updated version of DFS, Microsoft also extended the dfsrdiag.exe command line tool to include new functionalities.

The first of these extensions is the file hash function (DFSRDIAG.EXE FILEHASH). With this function you can compare the authoritative copy of a file against its replicated self by seeing if the file hashes are identical.

Furthermore, the new Replication State function (DFSDIAG.EXE REPLSTATE) allows you to analyze the current state of the replication service. With this, you can see what files are being updated on replication partners.

The basic idea behind another new function, ID Record (DFSRDIAG.EXE IDRECORD), is that every file and folder within a replicated folder has a corresponding ID record within the server's database, which is linked to valuable data like version and timestamp information.

With this function, you can determine a file or folder's record number and extract data bound to that record. This capability can be extremely helpful if you want to compare files stored on DFS replicas for consistency.

Overall, the changes Microsoft has made to DFS in Windows Server 2008 R2 should improve scalability and make the DFS easier to troubleshoot.

New use cases: File System Virtualization

Jonathan Brown, Publisher

While Microsoft has enhanced DFS to allow for larger configurations, the use-case for DFS is evolving. With many growing organizations, there is an increasing reliance on file virtualization as a means to reduce overall storage consumption and increase storage location flexibility. In this next article, author Brien Posey examines the pro's and con's of using DFS to create file system virtualization.

Add File Locking to DFSR to Prevent Version Conflicts Across Mirrored Servers

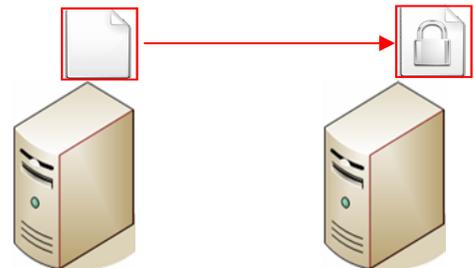
As companies expand, IT administrators are faced with the challenge to ensure that access to information is readily available to every employee, regardless of whether the person works in the headquarter building or a branch office location. This challenge is increased if the data files are large and the team members are geographically separated. To solve this problem, companies are relying on file mirroring solutions like Microsoft DFS Replication to ensure that a local copy of each file is instantly accessible from each office site, and that file changes are instantly replicated.

But when file mirroring solutions are used, how do you ensure that when one employee is updating a file another employee is not simultaneously making changes at another location and creating version conflicts?

PeerLock Server was designed to combat this issue by adding file locking technology to DFS Replication to provide a complete and efficient file collaboration environment. This assures that when a file is open at site A, all other versions - say local copies at various branch offices - are locked down, preventing anyone from opening and revising it. When the file closes, the file lock is released and the file is ready for synchronization.

Key features include:

- Application neutral file locking*
Lock files of any application type
- Cross domain locking support*
Allows for network logon credentials configuration
- Log file reporting*
See who opened what file and when
- Central monitoring console*
Review all current file use and file locking activity on the network
- Network disruption handling*
Through connection failure retry cycles



About Peer Software

Peer Software develops powerful, cost-effective, and easy-to-use data availability and collaboration software allowing organizations to seamlessly, efficiently, and reliably manage their digital assets. More than half of Fortune 100 companies entrust critical file management to Peer Software.

Download a free 15 day evaluation today at:

<http://www.dfsfilelocking.com/trial>

Using DFS to create file system virtualization in Windows Server 2008

Brien M. Posey, Contributor

As organizations accumulate more and more data, administrators often find that having a single file server does not meet their needs. Although adding additional file servers or even network-attached storage (NAS) devices is certainly an option, doing so introduces a degree of complexity into the organization because users' files are no longer being hosted in one single location.

In the past, an administrator had to keep track of which file resources were on which servers, and redirect users accordingly. This was confusing for both users and admins, especially when file resources had to be moved from one server to another to make up for disk space or other server resource shortages.

In Windows 2000 Server, Microsoft introduced a feature called the Distributed File System (DFS). This feature, which is also present in Windows Server 2003 and 2008, allows an administrator to virtualize the file system so the physical location of a user's files no longer matters. DFS creates a consolidated view of folders, even if those folders are spread across multiple file servers. Users access this consolidated view through a DFS root, which gives the illusion of being a normal share on a file server.

This form of file system virtualization is intended to help reduce confusion by presenting users with a unified view of the file systems of multiple file servers. Despite this, users still may not understand why they have to get to their files in a new way, or why they can suddenly see a bunch of files and folders that did not previously exist.

With those types of issues in mind, Microsoft has incorporated a few changes in Windows Server 2008 that make the file system virtualization process a bit less confusing for end users. Of these new features, the most helpful is Access-based Enumeration.

The basic idea behind Access-based Enumeration is that users are only allowed to see the files and folders they have access to. In previous versions of Windows Server, users accessing a virtualized file system through a DFS root could see the entire file structure, even if they didn't have permission to actually open all of the files and folders.

One thing to keep in mind about Access-based Enumeration is that it is not enabled by default (although it is enabled by default for newly created shares in Windows Server 2008). There are a couple of reasons for this. First, it makes the Windows Server 2008 DFS experience consistent with what users had when the DFS root was hosted on earlier versions of Windows Server.

A more important reason why Access-based Enumeration is not enabled by default, however, is that it's not compatible with earlier versions of Windows Server. If you want to use Access-based Enumeration, you must either be using a standalone DFS namespace that's hosted on a server running Windows 2008, or a domain based namespace that is running in Windows Server 2008 mode.

In case you have never heard of Windows Server 2008 mode, it is a new functional level offered by the Windows 2008 version of DFS. Setting Windows to use the Windows Server 2008 mode domain-based namespaces does two things:

1. Allows you to use Access-based Enumeration.
2. Increases the namespace's scalability.

DFS roots that are running the standard namespace mode (which is now called Windows 2000 Server mode) are limited to hosting about 5000 folders with targets. This is due to an Active Directory limitation found in older versions of Windows Server, which only allows a namespace directory object to grow to 5 MB. DFS roots running in Windows Server 2008 mode are not subject to this limitation. To the best of my knowledge, Microsoft has not released any information about what DFS is now capable of scaling to.

The requirements for using Windows Server 2008 mode are fairly straightforward. The domain must be set to use the Windows Server 2008 domain functional level, and all of your DFS namespace servers must be running Windows Server 2008.

Once you begin using Windows Server 2008 mode, you will have to manually implement Access-based Enumeration. You can do so by entering the following command:

```
DFSUTIL property abde enable \\<namespace_root>
```

You will also have to set the appropriate security permissions on each of the DFS folders, which you can do by using this command:

```
DFSUTIL property ACL grant  
\\<namespace_root>\<folder><username>:<permission>
```

Although implementing Access-based Enumeration is somewhat straightforward, it can be a bit complicated once you have DFS in place. Fortunately, Microsoft provides some really good resources that walk you through the process, including a basic overview of DFS in Windows 2008 and details Windows Server 2008 mode.

DFS – An Incomplete Solution

Jonathan Brown, Publisher

Microsoft has made some important strides with DFS for Windows Server 2008 including

- Failover clustering for greater scalability and availability
- Access rights based Enumeration
- File System Virtualization

But for organizations that have requirements that include collaboration on files between mirrored headquarter and branch office file servers being managed by DFS, there is a major shortcoming within DFS that companies should be aware:

- No File locking – lack of the ability to lock files in real time poses two (2) challenges
 - o Two colleagues can inadvertently edit the same file and the resultant replication will cause a version conflict
 - o Organizations using DFS for file server mirroring can expose the organization to version conflicts where all changes made prior to the last edit will be lost

Given this challenge, for organizations that are looking for the benefits of file replication, they should seriously consider third-party specialty products that address the functionality gap that is left by Microsoft. It is only through the addition of file locking technology that the goals of enhanced, secure, collaboration and unstructured information management can be met.

Resources from Peer Software



[PeerLock Adds File Locking to Microsoft DFS Replication to Prevent File Version Conflicts](#)

[PeerSync provides Multi-threaded Performance Alternative versus Microsoft DFS Replication](#)

[Business File Collaboration Increases Productivity and Efficiency for Team Projects](#)

About Peer Software:

Peer Software develops powerful, cost-effective, and easy-to-use backup and collaboration solutions allowing enterprises to efficiently manage their digital assets. The company's PeerSync software offers enterprise grade file synchronization, replication, distribution and backup capabilities. The PeerLock software offers distributed file locking to prevent version conflicts across mirrored file servers. Peer products are used by more than 8,000 enterprise customers world-wide, including over half of the Fortune 100.